

AFRL-IF-WP-TR-2005-1536

**GENETIC ALGORITHMS EVOLVE
OPTIMIZED TRANSFORMS FOR SIGNAL
PROCESSING APPLICATIONS**

**Frank Moore
Brendan Babb
Steven Becke
Heather Koyuk
Earl Lamson III
Christopher Wedge**

**University of Alaska, Anchorage
Department of Mathematical Sciences
CAS 154, 3211 Providence Drive
Anchorage, AK 99508**



APRIL 2005

Final Report for 10 January 2005 – 25 March 2005

Approved for public release; distribution is unlimited.

STINFO FINAL REPORT

**INFORMATION DIRECTORATE
AIR FORCE RESEARCH LABORATORY
AIR FORCE MATERIEL COMMAND
WRIGHT-PATTERSON AIR FORCE BASE, OH 45433-7334**

NOTICE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory Wright Site Public Affairs Office (AFRL/WS) and is releasable to the National Technical Information Service (NTIS). It will be available to the general public, including foreign nationals.

THIS TECHNICAL REPORT IS APPROVED FOR PUBLICATION.

/s/

PAT MARSHALL

Project Engineer

Embedded Information Systems Branch

/s/

AL SCARPELLI

Supervisor

Embedded Information Systems Branch

/s/

JAMES WILLIAMSON

Chief

Embedded Information Systems Branch

Advanced Computing Division

This report is published in the interest of scientific and technical information exchange and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>						
1. REPORT DATE (DD-MM-YY) April 2005		2. REPORT TYPE Final		3. DATES COVERED (From - To) 01/10/2005 – 03/25/2005		
4. TITLE AND SUBTITLE GENETIC ALGORITHMS EVOLVE OPTIMIZED TRANSFORMS FOR SIGNAL PROCESSING APPLICATIONS				5a. CONTRACT NUMBER F30602-03-2-0081		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER GAWAVE05IH		
6. AUTHOR(S) Frank Moore Brendan Babb Steven Becke Heather Koyuk Earl Lamson III Christopher Wedge				5d. PROJECT NUMBER SUNY		
				5e. TASK NUMBER 03		
				5f. WORK UNIT NUMBER IT		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) University of Alaska, Anchorage Department of Mathematical Sciences CAS 154, 3211 Providence Drive Anchorage, AK 99508				8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Information Directorate Air Force Research Laboratory Air Force Materiel Command Wright-Patterson AFB, OH 45433-7334				10. SPONSORING/MONITORING AGENCY ACRONYM(S) AFRL/IFTA		
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER(S) AFRL-IF-WP-TR-2005-1536		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.						
13. SUPPLEMENTARY NOTES Report contains color.						
14. ABSTRACT This report describes a genetic algorithm that evolves optimized sets of coefficients for signal reconstruction under lossy conditions due to quantization. The primary goal of the research described in this final report was to establish a methodology for using genetic algorithms to evolve coefficient sets describing inverse transforms and matched forward/inverse transform pairs that consistently outperform wavelets for image compression and reconstruction applications under conditions subject to quantization error.						
15. SUBJECT TERMS Genetic algorithms, discrete wavelet algorithms, quantization noise						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT: SAR	18. NUMBER OF PAGES 50	19a. NAME OF RESPONSIBLE PERSON (Monitor) Pat Marshall 19b. TELEPHONE NUMBER (Include Area Code) (937) 255-6548 x3586	
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified				

Abstract. The primary goal of the research described in this final report was to establish a methodology for using genetic algorithms to evolve coefficient sets describing inverse transforms and matched forward/inverse transform pairs that consistently outperform wavelets for image compression and reconstruction applications under conditions subject to quantization error. This report describes each of the following outcomes:

1. Coefficients trained on a single representative image consistently outperformed wavelets when subsequently tested against other images.
2. The performance of transforms evolved against representative subimages approximated that of transforms trained on the entire parent image(s), and was consistently better than the performance of the corresponding wavelet.
3. This research investigated the relationship between a subimage's representativeness and the performance of coefficients evolved against that subimage during subsequent testing.
4. This research extended the genetic algorithm to evolve coefficient sets describing matched forward and inverse transform pairs that further reduced mean squared error in quantized, reconstructed images, and identified a Pareto optimal front representing the tradeoff between compressed file size (FS) and maximum error reduction.
5. Attempts to use a genetic algorithm to evolve sets of coefficients that outperformed the standard 9/7 wavelet used by Joint Photographic Experts Group (JPEG) 2000 for image reconstruction were unsuccessful.

Outcomes 1-4 strongly support continued research into this exciting area.

1. SUMMARY

Wavelets have become the standard methodology for high-fidelity signal compression and reconstruction. Unfortunately, the performance of wavelets degrades under conditions in which the source signal is subjected to a significant amount of quantization error and/or noise. The primary goal of this research was to establish a methodology for using genetic algorithms (GAs) to evolve coefficient sets describing inverse transforms and matched forward/inverse transform pairs that consistently outperform wavelets for image compression and reconstruction applications under conditions subject to quantization error.

This report summarizes the results of research carried out by undergraduate students Brendan Babb, Steven Becke, Heather Koyuk, Earl Lamson III, and Christopher Wedge at the University of Alaska Anchorage (UAA) during the spring 2005 semester. This research was performed under the supervision of Dr. Frank Moore, the principal investigator (PI), who is an assistant professor of computer science in the UAA Mathematical Sciences Department. Funding for Dr. Moore's students was provided by the Research Foundation of the State University of New York, under contract to the Air Force Research Laboratory (AFRL).

Five primary subtasks were identified for this research. The goal of subtask 1 was to determine whether an inverse transform evolved using a single representative image also exhibited superior performance (in comparison to a selected standard wavelet) when subsequently tested against images other than the training image. The results summarized in section 3.1 of this report clearly substantiate the generalization properties of the inverse transforms evolved during this research. In particular, coefficients trained on a single representative image consistently outperformed wavelets when subsequently tested against all other images in the test set.

The goal of subtask 2 was to determine whether one or more representative subimages could be used to evolve an inverse transform whose performance, as measured by mean squared error (MSE) in the reconstructed image, consistently improved upon that of the corresponding wavelet inverse transform. The results summarized in section 3.2 clearly demonstrate that coefficient sets describing inverse transforms evolved against a training set of representative subimages consistently outperformed the corresponding wavelet. Further, these transforms performed approximately as well as transforms evolved against the entire parent image or images. By greatly reducing the amount of computation necessary to evaluate the fitness of each evolved transform, the use of subimages (rather than larger images) in the training population accelerated the evolutionary process by two orders of magnitude.

The use of subimages to evolve coefficients raised an interesting issue: what causal relationship exists between the representativeness of the subimages selected for training and subsequent performance of the evolved transforms on larger images? The goal of task 3 was to investigate this relationship. This research conclusively proved that use of more representative subimages during training will generally result in coefficients that exhibit better performance when reconstructing larger images. Further, these results suggest that factors such as the clarity and texture of the subimages play an important role: in particular, clear subimages with higher energy content consistently produce better evolved inverse transforms than blurry, dull subimages.

The goal of task 4 was to evolve coefficient sets describing matched forward and inverse transform pairs that further reduced MSE in quantized, reconstructed images. The earliest attempts to solve this problem revealed an interesting phenomenon: the GA automatically learned to boost each coefficient from the forward transform by a factor that was sufficiently large enough to offset much of the destructive effect of subsequent quantization. This phenomenon greatly improved the quality of the resulting reconstructed images, but at the cost of much

degraded file compression capability. An effective solution to this problem necessitated the identification of the pareto-optimal front representing the tradeoff between compressed FS and maximum MSE reduction. On one extreme of this front, the GA evolved coefficient sets that produced images of equal fidelity as those produced by the wavelet transform, but with much smaller compressed FS. As the FS constraint was relaxed, the GA evolved coefficients capable of much higher fidelity image compression and reconstruction. The result of subtask 4 was to demonstrate that the simultaneous evolution of matched forward and inverse transform pairs produced much lower MSE values than either the standard wavelet forward and inverse transform pairs or the combination of a wavelet forward transform and an evolved inverse transform.

The goal of subtask 5 was to incorporate a GA into a widely used wavelet analysis package to evolve sets of coefficients for forward and inverse transforms that outperformed the standard 9/7 wavelet used by JPEG 2000 for image compression and reconstruction. As with the first four subtasks, the GA seeded the initial population with randomly perturbed copies of the standard wavelet (in this case, the 9/7 wavelet), and used the GA in an attempt to evolve new coefficients for transforms having an identical structure as the wavelet. To date, none of the attempts to produce a GA capable of consistently evolving such coefficients have been successful. It is possible that, by incorporating the subimaging methodology of subtasks 2 and 3 and using significantly larger scale runs, the GA will ultimately be able to improve upon the compression and reconstruction capabilities of the JPEG 2000 standard.

The outcomes of subtasks 1, 2, 3, and (especially) 4 strongly support continued research into this exciting area. Integration of the subimage training methodology from subtask 2 with the simultaneous evolution of forward and inverse transform pairs from subtask 4 should allow us to make considerably more rapid progress toward the goal of revolutionizing the field of image compression and reconstruction under conditions subject to large quantization error. Additional research into the issue of subimage representativeness raised during subtask 3, and its impact upon the quality of reconstructed images, may allow us to identify techniques for identifying, highlighting, and extracting highly distinctive subimages (e.g., a tank moving across the desert) within larger images. Furthermore, the identification of the properties of subimages that make them more or less suitable for GA training will allow us to automatically enhance the precision of the subimage-based approach.

2. INTRODUCTION

Wavelets have become the standard methodology for high-fidelity signal compression and reconstruction. For example, the JPEG 2000 image coding system is based upon wavelet technology, and is used for applications ranging from medical imaging to portable digital cameras. Unfortunately, the performance of wavelets degrades under conditions in which the source signal is subjected to a significant amount of quantization error and/or noise. State-of-the-art wavelet-based signal compression techniques, thus, leave considerable room for improvement.

The payoff from improving upon standard wavelet technology could be extraordinary. Better medical images would mean that doctors could more consistently and more accurately identify cancer early in its development. Better satellite images would result in more accurate identification and tracking of objects of interest. Better audio compression and reconstruction techniques would mean that sounds produced by telephone receivers, CD and DVD players, and other devices would be more realistic and more understandable. Images downloaded from the internet would be smaller, more quickly transmitted, and/or more accurate. The quality of digital photography would improve. These are but a few of the many endeavors that would benefit from improved signal and image processing technologies.

To date, there have been and continue to be many attempts to identify new wavelet transforms and inverse transforms that improve upon the state of the art for a wide variety of application areas. However, it is not known whether any of the researchers involved in these attempts has considered the possibility that the use of nontraditional transforms—inspired by wavelets but not necessarily bound by their precise mathematical properties—could actually result in improved signal and image reconstruction. Hence, there is virtually no literature directly related to the proposed project, apart from two final reports produced by the PI summarizing the results of his prior work in collaboration with the Air Force.

Previous research conducted by the PI established a methodology for using a GA to evolve nontraditional inverse transforms for signal reconstruction applications. The evolved inverse transforms consistently outperformed wavelets: in particular, these transforms reduced the MSE observed in specific classes of reconstructed one-dimensional (1-D) signals by a factor of 10 or more, in comparison to the standard (and widely used) Daubechies-4 (Daub4) inverse wavelet transform. Furthermore, this research demonstrated that a reduction in MSE was also possible for multidimensional signals, such as two-dimensional (2-D) images. These results laid the foundation for extensive additional research in this exciting new area.

Previous research was accomplished through the development of two specialized GA packages. The first package was designed to rapidly evolve a novel inverse transform for 1-D signals, such as sine waves or ramp functions. This package can efficiently process populations of thousands of candidate solutions over thousands of generations. The second GA package was designed to optimize an inverse transform for reconstructing 2-D images. Together, these two GAs provided a solid foundation upon which the software developed for the proposed research could be built.

The PI's prior research was restricted to the evolution of coefficients for inverse transforms. Prior to this investigation, no research had been conducted to determine whether the coefficients for a forward wavelet transform could be used as a starting point for evolving a forward transform that reduces MSE due to quantization or noise. Likewise, no research had attempted to evolve matched forward and inverse transform pairs; such matched transforms may result in even greater error reduction than is possible from a single evolved transform. Finally, it was not known whether a GA could evolve optimized transforms whose structure and composition may radically

differ from that of any wavelet. The potential for high payoff strongly encourages pursuit of additional research in this critical technology area.

3 RESULTS AND DISCUSSION

There were five key questions to be answered during this project:

1. Could the coefficient sets trained on a single representative image perform consistently well when subsequently tested against other images?
2. How will the performance (in terms of MSE reduction) of transforms evolved against representative subimages compare to that of transforms trained on the entire parent image(s)?
3. What happens when a nonrepresentative subimage is used to evolve coefficients? What is the relationship between the representativeness of subimages used for training versus the effectiveness of the resulting coefficients when tested against the parent image or other images? Is it possible to evolve an inverse transform that is capable of highlighting each occurrence of a nonrepresentative subimage in a larger image?
4. Could the GA be modified to evolve coefficient sets describing matched forward and inverse transform pairs that further reduced MSE in quantized, reconstructed images? Forward transforms tend to compensate for compression error by exaggerating values from the original signal. Will it be possible to design the GA in a manner that allows the user to specify the maximum compressed FS produced by the forward transform?
5. Can the methodology established by this research identify new transforms to replace wavelets for a widely used application, such as the standard 9/7 wavelet used by JPEG 2000 for image reconstruction?

Each of these questions is directly addressed in the following subsections.

3.1 SUBTASK 1: Generalization Properties of Evolved Coefficients

Prior research conducted by the PI established two key properties of evolved inverse transforms:

- A GA is capable of evolving an inverse transform that reduces the MSE observed in reconstructed images previously subjected to quantization error.
- An inverse transform evolved against a training population consisting of several members of a given class of images is capable of high-quality reconstruction of test images subsequently drawn from the same class. In particular, the percentage reduction in MSE for images from the training and test sets were approximately equal.

The first subtask addressed by this research was to determine whether an inverse transform evolved against a single representative training image was capable of similar MSE reductions when subsequently tested against other images. The results of this subtask are summarized in Table 1. Each run evolved a population of $M = 500$ candidate solutions for $G = 200$ generations; these control parameters proved to be sufficient for the purposes of this subtask.

These results clearly show that each of the coefficient sets evolved using a single image from the training set outperforms the standard Daub4 wavelet, when tested against any other image from this set. Thus, all of the results from subtask 1 underscore the outstanding generalization capabilities of inverse transforms evolved against selected single images.

Brendan Babb and Steven Becke performed the data collection necessary to complete subtask 1.

Table 1. Generalization Properties of Inverse Transforms Evolved Against a Single Image

- a) Wavelet: Daub4
 Training Image: fruits
 Quantization: 64

Test Image	MSE (using Daub4)	MSE (evolved)	MSE as % of Daub4 MSE	Percentage Improvement of MSE
airplane	112.4558	107.98	96.01995	3.980052607
barb	307.6155	295.22	95.97046	4.029543375
fruits	115.8036	110.51	95.42881	4.57118777
goldhill	135.7847	129.09	95.06962	4.930378754
park	168.9597	162.35	96.088	3.91199795
susie	132.5596	127.41	96.11526	3.884743165

Averages 95.78202 4.217983937

- b) Wavelet: Daub4
 Training Image: couple
 Quantization: 64

Test Image	MSE (using Daub4)	MSE (evolved)	MSE as % of Daub4 MSE	Percentage Improvement of MSE
baboon	279.5612	262.56	93.91861	6.081387546
couple	155.9632	145	92.97065	7.029350513
fruits	115.8036	107.77	93.06274	6.937262745
lenna	177.5059	163.53	92.12652	7.873484769
peppers	121.2649	116.15	95.78204	4.217955897
zelda	134.5674	125.71	93.41787	6.582129104

Averages 93.5464 6.453595096

- c) Wavelet: Daub4
 Training Image: barb
 Quantization: 32

Test Image	MSE (using Daub4)	MSE (evolved)	MSE as % of Daub4 MSE	Percentage Improvement of MSE
airplane	43.07	41.19	95.63501	4.36498723
barb	96.89	92.91	95.89225	4.107751058
couple	59.67	56.81	95.20697	4.793028322
goldhill	55.81	53.5	95.86096	4.139043182
park	70.33	67.54	96.03299	3.967012655
susie	53.38	51.6	96.66542	3.334582241

Averages 95.88227 4.117734115

- d) Wavelet: Daub4
 Training Image: boat
 Quantization: 32

Test Image	MSE (using Daub4)	MSE (evolved)	MSE as % of Daub4 MSE	Percentage Improvement of MSE
baboon	113.77	109.67	96.39624	3.603761976
boat	52.52	49.92	95.0495	4.95049505
fruits	41.52	39.72	95.66474	4.335260116
lenna	63.55	59.71	93.95751	6.042486231
peppers	44.15	43.32	98.12005	1.8799547
zelda	34.52	32.83	95.10429	4.89571263

Averages 95.71539 4.284611784

- e) Wavelet: 2/6 (TS)
 Training Image: goldhill
 Quantization: 32

Test Image	MSE (using 2/6)	MSE (evolved)	MSE as % of 2/6 MSE	Percentage Improvement of MSE
airplane	40.72	39.89	97.96169	2.03831041
barb	93.73	92.89	99.10381	0.89619119
fruits	42.38	40.94	96.60217	3.39782916
goldhill	55.69	52.67	94.57712	5.42287664
park	68.49	66.78	97.50329	2.49671485
susie	54.13	52.22	96.47146	3.5285424

Averages 97.03659 2.96341078

- f) Wavelet: 2/6 (TS)
 Training Image: park
 Quantization: 32

Test Image	MSE (using 2/6)	MSE (evolved)	MSE as % of 2/6 MSE	Percentage Improvement of MSE
baboon	113	109.78	97.15044	2.84955752
couple	59.91	56.63	94.52512	5.47487899
park	68.49	65.52	95.6636	4.33639947
lenna	64.18	60.14	93.7052	6.29479589
peppers	45.02	42.54	94.49134	5.50866282
zelda	35.38	32.47	91.77501	8.22498587

Averages 94.55179 5.44821343

- g) Wavelet: 2/6 (TS)
 Training Image: susie
 Quantization: 64

Test Image	MSE (using 2/6)	MSE (evolved)	MSE as % of 2/6 MSE	Percentage Improvement of MSE
airplane	106.56	97.31	91.31944	8.68055556
barb	310	286.66	92.47097	7.52903226
couple	158.58	142.31	89.74019	10.2598058
goldhill	140.44	123.76	88.12304	11.8769581
park	167.17	151.61	90.69211	9.30789017
susie	133.23	118.91	89.25167	10.74833

Averages 90.26624 9.73376197

- h) Wavelet: 2/6 (TS)
 Training Image: peppers
 Quantization: 64

Test Image	MSE (using 2/6)	MSE (evolved)	MSE as % of 2/6 MSE	Percentage Improvement of MSE
baboon	277.97	261.81	94.18642	5.813577
boat	161.28	148.36	91.98909	8.0109127
fruits	117.17	107.53	91.77264	8.22736195
lenna	178.6	159.96	89.56327	10.4367301
peppers	124.52	112.27	90.16222	9.83777706
zelda	138.76	124.59	89.78812	10.2118766

Averages 91.24363 8.75637258

3.2 SUBTASK 2: Evolving Inverse Transform Coefficients Using Subimages

During the summer of 2004, the PI showed that a GA can be used to generate inverse transform coefficients which outperform standard wavelet inverse transforms for image reconstruction under conditions subject to quantization error. Unfortunately, computation time was prohibitive, taking as much as 46 hours to complete a single evolution run. The purpose of subtask 2 was to establish a methodology for evolving optimized inverse transform coefficients using a training population consisting of representative subimages. In addition, considerable effort was invested into improving the overall quality of the GA/wavelet software package.

Christopher Wedge completed subtask 2.

3.2.1 Representative Subimages

Subtask 2 began with the problem of choosing representative subimages for training. This process seemed to be highly subjective. There were no criteria listed in previous results that could be used to determine whether a given image was representative of the parent image. The approach adopted for subtask 2 was to use subimages that were duplicates of the original, but were substantially reduced in size.

A total of 63 tests were performed. For the first 60 tests, the GA used a single 32- by 32-pixel subimage (LennaMini.bmp, MonetWaterLilliesMini.bmp, BarbMini.bmp, and GoldhillMini.bmp, each) to evolve optimized inverse transform coefficients. The final three tests used all four of these subimages.

Each of the 60 single-subimage tests were repeated 5 times with quantization levels of 32 and 64, for a total of 10 training runs on each individual image. The final 3 tests were run once each with quantization levels of 32 and 64, for a total of 2 training runs per test. Each training run used values of (population size) $M = 5,000$, and (number of generations $G = 2,500$). (Note that M is 10 times and G is 5 times those used in the 46-hour runs completed by the PI during the previous investigation.)

Multiresolution (MR) analysis was neglected for time reasons (exploring MR levels of 2 and 3 would double and triple the time to run, respectively, and would triple the number of tests) and because it was thought that any change in MSE would likely be exacerbated by using different MR values. That is, any improvement with $MR = 1$ would still be present, if not amplified, with higher values. Threshold was similarly ignored for time reasons, and set to 0. Quantization tends to have a bigger effect on MSE values unless the threshold is set exceptionally high.

In each run, the initial coefficients used were taken from the Daub4 wavelet, for both time reasons, and because the original research also focused on the Daub4 wavelet.

The results of these tests are represented by Tables 2 and 3. On average, the training runs using a single subimage were approximately 50 times faster than the corresponding runs performed on the entire parent image. In every instance, the evolved transform coefficients outperformed the Daub4 wavelet when reconstructing the original full-sized parent image. Further, the evolved transforms outperformed the Daub4 wavelet when reconstructing other full-sized images (e.g., a transform evolved on the miniature version of lenna.bmp also outperformed Daub4 on barb.bmp, goldhill.bmp, and MonetWaterLillies.bmp) in all tests. Mean MSE reduction over Daub4 wavelet on lenna.bmp was 2.41 percent and 8.32 percent with quantization of 32 and 64, respectively. Mean MSE reduction on barb.bmp was 2.60 percent and 5.67 percent with quantization of 32 and 64, respectively. Mean MSE reduction on MonetWaterLillies.bmp was 3.67 percent and 5.55

percent with quantization of 32, and 64, respectively. Mean MSE reduction on goldhill.bmp was 2.45 percent and 8.33 percent with quantization of 32 and 64, respectively.

Table 2. Generalization Properties of Coefficients Evolved Using Subimages
(Q = quantization)

Trained Image	Q	Evolution Time		Performance vs. Lenna		Performance vs. Barb		Performance vs. Monet		Performance vs. Goldhill	
		Mean	St Dev	Mean MSE	St Dev	Mean MSE	St Dev	Mean MSE	St Dev	Mean MSE	St Dev
Lenna Mini	32	1:01:38	0.0007	62.36	0.171	94.50	0.119	78.69	0.143	54.55	0.138
	64	1:02:33	0.0003	163.56	0.426	291.54	0.360	241.26	0.444	125.66	0.313
Barb Mini	32	1:02:02	0.0005	62.49	0.633	94.18	0.399	78.15	0.402	54.38	0.392
	64	1:01:26	0.0008	164.82	0.324	289.27	0.249	242.36	0.437	125.41	0.364
Goldhill Mini	32	1:01:05	0.0008	62.73	0.277	95.73	0.235	78.96	0.263	55.61	0.208
	64	1:01:37	0.0005	162.87	0.511	291.10	0.436	241.84	0.565	124.92	0.377
Monet Mini	32	1:01:35	0.0006	60.47	0.229	93.06	0.190	77.70	0.218	53.25	0.131
	64	1:01:47	0.0005	159.70	0.493	288.80	1.578	242.50	0.370	121.89	0.258
All Singles	32	1:01:35	0.0006	62.01	0.985	94.37	1.007	78.38	0.558	54.45	0.886
	64	1:01:51	0.0006	162.74	1.983	290.17	1.426	241.99	0.654	124.47	1.581
DAUB4	32	N/A	N/A	63.55	N/A	96.89	N/A	81.36	N/A	55.81	N/A
	64	N/A	N/A	177.51	N/A	307.62	N/A	256.21	N/A	135.78	N/A
The Four Minis	32	4:13:27	N/A	60.48	N/A	93.13	N/A	77.14	N/A	53.33	N/A
	64	4:07:33	N/A	160.89	N/A	287.42	N/A	239.82	N/A	122.84	N/A

Table 3. MSE Reduction of Coefficients Evolved Using Subimages
(Relative to the Daub4 Wavelet)

Trained Image	Quantization	Mean MSE Reduction Over Daub4			
		vs. Lenna	vs. Barb	vs. Monet	vs. Goldhill
Lenna Mini	32	1.865531%	2.465221%	3.276650%	2.267041%
	64	7.855515%	5.226341%	5.834097%	7.456592%
Barb Mini	32	1.655071%	2.789749%	3.276650%	2.578624%
	64	7.145221%	5.964889%	5.834097%	7.640751%
Goldhill Mini	32	1.281025%	1.189940%	3.941018%	0.370489%
	64	8.246600%	5.369706%	5.406229%	7.998091%
Monet Mini	32	4.835340%	3.954606%	4.497169%	4.594401%
	64	10.034279%	6.117590%	5.349971%	10.231439%
All Singles	32	2.409242%	2.599879%	3.666830%	2.452639%
	64	8.320404%	5.669631%	5.549295%	8.331718%
The Four Minis	32	4.822056%	3.872319%	5.183125%	4.451723%
	64	9.362846%	6.566303%	6.397766%	9.533019%

As expected, the runs which trained on all four miniature images generally outperformed the single-image runs. However, additional testing will be necessary to conclusively prove the benefits of using multiple subimages to evolve inverse transform coefficients.

Still more testing should be done with varying levels of threshold, MR analysis, differing base wavelets, multiple image evolution, etc. Nevertheless, subtask 2 demonstrated that there is performance to be gained, in terms of computation time and MSE reduction, by genetically evolving wavelet transform coefficients on representative subimages (in this case, miniature versions of the original image).

3.2.2 Nonrepresentative Subimages

An interesting phenomenon noticed by Tinsley and Kettle was that transform coefficients trained on a nonrepresentative subimage performed better than the Daub4 wavelet on that subimage, but worse than the Daub4 wavelet on the entire parent image. It stood to reason that coefficients evolved on nonrepresentative subimages might feasibly be used to efficiently search an image for an instance of particular subimage, effectively highlighting the subimage and subduing the remainder of the parent image.

Eighteen tests were run by co-evolving a nonrepresentative subimage (in each case, the subimage was LennaEye.bmp) against the parent image (in each case, the super image was lenna.bmp) and using the GA to attempt to simultaneously maximize the MSE of the reconstructed parent image and minimize the MSE of the reconstructed subimage.

The runs were broken down into three categories, based on weighting these two fitness criteria. The three categories, in terms of subimage weight to parent image weight were: 1 to 1, 2 to 1, and 1 to 2. The six runs in each of the three categories consisted of tests of every combination with quantization of 0, 32, and 64, and threshold of 0 and 16. MR levels were kept at 1, and the base wavelet used was again Daub4.

Since the co-evolution involved transforming both the 32- by 32-pixel subimage and the 512- by 512-pixel parent image, M and G were chosen to be 500 and 200, respectively. Even with such small run parameters, the mean runtime was 4 hours, 2 minutes, 58 seconds.

Table 4 tabulates the results of these tests. Lenna.bmp was used as the parent image and Lenna's eyeball was used as the subimage. Evolved transforms were then iteratively (25 and 50 times) applied to the parent image to determine if the subimage was being highlighted. Within the iterative transform application loop, various methods were attempted. One approach was to simply apply the forward transform, quantization, and inverse transform for the desired number of iterations. However, it was found that if a quantization value above 0 was used, then iterative applications produced no effect after the first. It seems the error introduced by applying the transform more than once was being completely offset after quantization, and that the residual transform was subsequently able to perform lossless compression in the presence of the quantization. Another approach was to transform only the Y values, leaving U and V alone. Yet another involved repeatedly applying the forward transforms, then quantizing, and finally applying the same number of reverse transforms.

Each approach was followed by the user judging the resulting image to determine if the subimage was effectively highlighted. This is certainly not a quantitative process, and relies entirely on the user's opinion, but since the goal was to produce an image from which humans can easily distinguish a desired subimage, it was unavoidable. As a result, there is a lot of subjectivity as to

whether or not a subimage has been effectively highlighted. However, there can be no doubt that the evolved transforms did not produce anything near the desired effect in these tests.

Often the application of the Daub4 base transform to both the subimage and parent image produced MSE values which were substantially higher on the subimage than on the parent image. While the GA did start to produce better transforms, the low values of M and G did not allow for much time for evolution to work its magic. Further, it is speculated that perhaps due to the relatively limited search space of the GA, the evolved coefficients are simply too close to the mathematically ideal wavelets to solve this problem more effectively. That is, the initial transforms produce positive results universally, and to see the large disparity needed between subimage and parent image MSE values to highlight a subimage, the search space needs to be broadened. Unfortunately, such an algorithm would have to search far more to produce good results, which necessitates far higher values of M and G, naturally costing far more time and computing power to run.

Table 4. Results of Attempting to Evolve Inverse Transforms that Perform Well Only on Selected Subimages

Q	Threshold	Fitness Weights (sub vs. super)	Run Time	MSE on Sub	MSE on Super	D4 on Sub	D4 on Super
0	0	1 to 1	4:02:53	0.547851563	0.67231369	0.737630208	0.706765493
0	16	1 to 1	4:02:21	31.45572917	26.02457937	33.53515625	18.37425359
32	0	1 to 1	3:59:53	125.1875	99.44009781	107.4007161	63.54538091
32	16	1 to 1	4:08:27	122.2522786	97.34475835	107.4007161	63.54538091
64	0	1 to 1	4:00:50	284.7936198	250.6794001	256.733724	177.5058975
64	16	1 to 1	4:00:30	280.4238281	243.8678487	256.733724	177.5058975
0	16	2 to 1	4:02:48	31.81575521	17.16363017	33.53515625	18.37425359
32	0	2 to 1	4:03:03	110.9404297	76.9559199	107.4007161	63.54538091
32	16	2 to 1	4:03:29	109.1940104	75.45477931	107.4007161	63.54538091
0	0	1 to 2	4:01:39	0.688802083	0.76631546	0.737630208	0.706765493
0	16	1 to 2	4:02:22	31.81575521	17.16363017	33.53515625	18.37425359
32	0	1 to 2	4:02:08	143.9163411	129.7384529	107.4007161	63.54538091
64	0	1 to 2	4:01:03	330.9960938	318.8694585	256.733724	177.5058975
Average Run Time			4:02:25				
StDev			0.001466				

3.2.3 Software Enhancements

Major additions and bug fixes are listed. Minor changes, fixes (of which there are a great many) are not.

- Separated transform evolution and application: Before they were one process, now they are individual, unrelated processes.
- Redesigned how transform coefficients are stored: They are now persistent. Before, only the inverse transform coefficients were stored, now both forward and inverse are. This feature is nice for extending the subimage training software to simultaneously evolve forward and inverse transform coefficients. This modification required sweeping changes to much of the

previous code base, including most of the methods in GAWavelet.cpp and WaveletSettings.cpp.

- Added the ability to save transform coefficients to a file: Previously, the user had to enter debug mode to harvest the coefficients.
- Wrote a small program to read the coefficients from a saved transform file.
- Added the ability to load transform coefficients from a previously saved file: This technique allows for much easier test result reproduction and is a logical analog to the save feature. Also, by loading previously evolved coefficients, the user can then further evolve them in a new run.
- Added a separate GA which co-evolves a subimage with its parent image, trying to maximize MSE of transform application to the parent while minimizing the MSE of application to the subimage.
- Added a search image feature. As yet, the search image logic (which has changed frequently over the course of testing) has yet to accomplish the desired task. It is still a work in progress.
- Fixed a serious bug in which the MSE value given on any transform applied after the first was incorrect.

3.2.4 Recommended Future Software Modifications

- Debugging: There are still known bugs present in the program. These include out of memory errors when the user manually types values into the wavelet settings window. When opening, then saving bitmap files, two bytes are appended, and the internal file structure is vastly changed. There are also likely undiscovered bugs lurking about.
- Make the application operating system (O/S) independent: It would be nice to not be so dependent on Windows.
- Rewrite the application with parallel computing in mind: The smallest runs computed took over an hour, with larger runs taking far longer. There is much to be gained from a parallel processor environment.
- Overhaul the application: The application used as a starting point for this research was poorly designed and buggy from the start. As each person has added their own features, the program has become more and more obfuscated. It would be worthwhile to completely redesign and rewrite the existing program with future extension in mind.

3.2.5 Recommended Future Research Directions for Subimage Processing

- Variable length transforms: Currently, the application evolves coefficients which adhere to the existing length of the base wavelet (e.g., the Daub4 wavelet transforms each have four coefficients—evolving off the Daub4 wavelet results in transforms which also have four coefficients). It is possible that exploring variable length transforms may further improve results. Modifications to the way in which the application stores, saves, and loads transforms should easily allow for this extension.
- Integration of simultaneous evolution of forward transform coefficients with subimage training.
- Co-evolution of U and V values: Currently only the Y values are used in the GA. Although the majority of the picture information is stored in the Y domain, expanding to also incorporate U and V may be fruitful.
- Expand the search space: Presently the GA only searches in the transform space immediately adjacent to the base wavelet. Expanding the search could have very big ramifications, especially for the subimage enhancement and extraction problem, since one of the objectives of that problem is to maximize the MSE of the parent image.
- Improve the GA: The GA, as it stands, is fairly simple. Including more advanced techniques such as niching, linear normalization, etc. could generate better results.

- Explore other wavelets: Only the Daub4 wavelet was used during subtask 3 due to time constraints. Further exploration using the other wavelets could be fruitful.
- Test other quantization levels: These results indicate that the advantage over standard wavelets increases proportionately with quantization error. Unfortunately, only two nonzero quantization levels were used during this subtask. Tests utilizing additional levels of quantization error will be necessary to conclusively demonstrate the relationship between quantization error and the percentage improvement in MSE obtained via the GA.

3.3 SUBTASK 3: Subimage Representativeness and Reconstructed Image Quality

The objectives of this subtask were a) to determine how important representativeness is to GAs trained on selected images or subimages, and b) to experiment with how representativeness could be determined and measured: could a formula be discovered that would adequately distinguish representative from nonrepresentative images?

Heather Koyuk completed subtask 3.

3.3.1 The Issue of Representativeness

Although reason indicates that the more representative an image is of another, the better a transform evolved on one will perform on the other, some sort of substantiation is necessary. The questions of how this representativeness can be measured, how important it actually is to the performance of an evolved GA, and how an evolved transform might be able to be trained to minimize (or, in some cases, maximize) the differences between images are all important to this project as a whole.

3.3.2 Evolving Inverse Transforms on Subimages

By using smaller subimages to evolve inverse transforms, the amount of time it takes to come up with notable results can be significantly decreased with little effect on results. Although work has been done on this before, further testing and validation of the results appeared necessary.

3.3.3 Design of the Investigation

A number of inverse transforms were evolved using representative and nonrepresentative subimages. These GAwavelets were then tested on each of the larger images and the results were analyzed in an attempt to determine how the originating subimage influenced the result on the larger images. At the same time, the results were examined for evidence that the subimage-trained transforms could decrease the MSE relative to standard wavelet-based techniques.

3.3.4 Representativeness Formula

The formula used to determine representativeness was the following, where stdev. means the standard deviation and Y, U, and V stand for the Y, U, and V values of the image respectively:

$$\sqrt{(Stdev(Y))^2 + ((Stdev(U))^2)/4 + ((Stdev(V))^2)/4} . \quad (1)$$

More weight was given to the Y values for a number of reasons:

- The current GA evolves coefficients based on the Y values only.
- One of the main advantages of YUV format is that chrominance (which consists of the U and V values) is much less important to data reconstruction than luminance (Y) is.
- Except for the two color images used (barb.bmp and lenna.bmp), U and V values were basically nonexistent, with an average mean of -0.0165 and 0.0001, respectively, and an average standard deviation of 0.1125 and 0.1083, respectively (as opposed the average Y

mean of 128.8 and Y standard deviation of 48.7). Barb.bmp and lenna.bmp had an average U and V means of -10.5 and 27.6 and average U and V standard deviations of 12.3 and 13.5).

3.3.5 Procedures

After saving 24 subimages produced by the GA Java code, the C++ code obtained from the PI was executed to train the coefficients for inverse transforms against each of the subimages. Due to time constraints as well as a desire to produce results comparable with those of Tinsley and Kettell, this experiment used the same settings that Tinsley and Kettell used: quantization of 64, population size $M = 500$, and maximum number of generations $G = 500$. After saving the evolved coefficients, each of the coefficient sets were tested against each of the 12 larger original images. Results were compiled into a number of data tables. The results were then analyzed for patterns of behavior.

3.3.6 Representativeness Results and Confirmation of Tinsley's and Kettell's Results

Not surprisingly, the coefficients evolved on the more representative subimages performed significantly better than the coefficients evolved on the nonrepresentative subimages. In fact, the nonrepresentative coefficients had an average percentage increase in MSE (relative to the Daub4 wavelet) of 4.86 percent, whereas the representative coefficients had an average percentage reduction in MSE of 5.43 percent, for a greater than 10 percent difference between the two groups. Perhaps more significantly, the nonrepresentative coefficients had an average standard deviation (among the individual images) of 28, whereas the representative coefficients had an average standard deviation of 2.9. The averaged variance for the nonrepresentative coefficients was 1445, while the representative coefficients had an average variance of only 9 (see table 4, Appendix). This result preliminarily indicated that the standard deviation score initially used for this subtask produced a relatively reasonable measure of representativeness. This result also further substantiated Tinsley and Kettell's preliminary findings, indicating that subimages could be used to train GAs, provided that the subimage used confirmed to certain standards.

3.3.7 Overall Image Results

What was unexpected was that the GAs generally performed well over the whole group of images: instead of diagonal effect where coefficients trained on subimages of an image performed better only on that particular image, the effect was rather linear, as shown in Figure 1 and 2. Upon examining the bitmap files that bred each of these coefficient sets (shown in Figure 3), a pattern began to emerge. Although specifically difficult to distinguish individually by the eye (e.g., what makes the GA from fruits_least better than the one from Susie_least), the images patently progress from clear, textured images to flat, blurry ones. Obviously, this is an important detail that should be further explored while working with the evolution of transform coefficients. This would also indicate that representativeness of the training image or subimage does not play as large a factor in coefficient success as was first assumed, and that some factor for clearness and/or texture is more important. Further research should be conducted in order to clearly identify and quantify this factor.

3.3.8 Influence of Image Used for Testing

Certain images also tended to respond better to the GAs overall (see Figure 3), but it was not as easy to distinguish what made those images easier for the GAs, due to their size and complexity. It may be worth investigating whether wavelet coefficients can be evolved that can reduce the MSE on the bad files.

3.3.9 Conclusions, Subtask 3

Although the initial results for subtask 3 seem to indicate that representativeness is not as important a factor as initially thought, these tests did reveal that certain subimages might provide

better results than others. Such an effect should be explored further. Many more tests should be run to confirm these initial findings, with varying factors such as changing the fitness function in the Java subimage selection code, testing the evolved coefficient sets on completely different images, and running the GA for longer lengths of time. These tests did reconfirm Tinsley and Kettell's previous results with the evolution of coefficients on subimages of smaller size, provided that the subimage used to train the GA confirmed to certain standards.




Figure 1. Percentage Improvement by Image




Figure 2. Percentage Improvement by Subimage



Figure 3. Subimages Ordered by the success of Their Evolved GA (left to right, top to bottom)

3.4 SUBTASK 4: Evolving Matched Forward and Inverse Transform Pairs

The purpose of subtask 4 was to establish a methodology for evolving matched forward and inverse transform pairs for improved image compression and reconstruction.

Brendan Babb and Steven Becke completed subtask 4.

3.4.1 Discussion

The first step in this process was to modify software previously developed by the PI to simultaneously evolve coefficients for both the forward and inverse transforms. Initial attempts to minimize the MSE produced by the evolved transforms produced good results, improving the MSE in comparison to the performance of the standard D4 wavelet by over 30 percent. Taking the coefficients evolved during one run and plugging them back into the program as the starting coefficients resulted in up to a 95 percent reduction in MSE compared with the standard wavelet MSE.

At this point it became evident that something must be wrong because these initial results were simply too good. The original fitness function focused entirely on improving MSE. Since the new approach simultaneously evolved both the forward and inverse coefficients, it soon became obvious that the GA was merely evolving a forward transform that increased the FS in order to offset the effects of the quantization step: in effect, the GA learned to make the Y values larger so that all the detail would be retained when the Y values were quantized.

To correct this error, both FS and MSE were integrated into the fitness function used by the GA. The tradeoff between these two conflicting objectives defines in a Pareto optimal front that must be handled correctly by the fitness function. Since the GA was storing the Y values after quantization for the compressed file, the fitness function was modified to compute the sum of the log base 2 of the Y values, as this should correspond to the number of bits to store each Y value. To avoid taking the log of 0, this code was modified as follows:

```
if (Y == 0)
    Y length = Y length + 1;
else if (Y > 1)
    Y length = Y length + log2(Y);
```

This approach proved to be erroneous due to the fact that there were actually negative Y values. Thus, modifications were made to include bits for sign and also to take the absolute value of Y. This led to improved coefficients, but when the coefficients were plugged back into the original GA, the expected FS deviated from the actual FS.

The final modification was to include a call to Encode Frame. This call was avoided initially as it was thought to add too much time to the GA to prepare a buffer for saving the file each time. When it was added it did give accurate FS calculations, and also slowed down the speed of running the GA. Currently, the GA only does an Encode Frame operation on the Y component, so that the function only has to be called once in the GA. The Encode Frame of U and V are assumed to be the same each time and add 104 bytes to the actual FS.

Testing revealed a discrepancy in the MSE improvements based on the GA and the actual MSE. This error can be attributed to the fact the GA minimizes MSE for the Y values and not the final red/green/black (RGB) values. The difference is only 2 to 3 percent lower in the MSE improvement. This error could be rectified in future versions by converting back to RGB and

computing MSE, but this solution would slow down the GA and a mathematical approximation could probably be found.

Another possible way to speed up the program would be to use a mathematical approximation for Encode Frame that would give an accurate FS.

The initial fitness function was using of the form $a * \text{MSE ratio} + b * \text{FS ratio}$. It was challenging to pick the right combination of a and b . Depending on their values, the evolved coefficients might have a good MSE gain but also grow in FS. To compensate for this effect, the GA was modified to control the maximum FS allowed during each run and see what improvements in MSE could be obtained for the specified maximum FS.

The image Couple, the D4 wavelet coefficients, and a quantization step of 64 were used for most of the tests performed during this research. Trail-and-error revealed two important metrics:

- $\text{MSE Ratio} = 100 * \text{MSE} / (\text{the original MSE for couple with quantization 64 and the D4 coefficients}).$
- $\text{FS Ratio} = 100 * \text{FS} / (\text{original FS for couple with quantization 64 and the D4 coefficients}).$

Many combinations of $a * \text{MSE ratio} + b * \text{FS ratio}$ were tried as a measure of fitness. Coefficients evolved during one run of the GA could be used as a starting point for a subsequent run, i.e., it was possible to plug these coefficients back into the program, change a and b slightly, and get better results. It was almost a tuning of sorts: one could start with a solution that was approximately 100 percent FS ratio and 85 percent MSE ratio, modify the a and b values, rerun the program, and end up with a solution with approximately a 70 percent MSE ratio but 108 percent FS ratio.

In order to look for a particular FS percentage, the fitness function was altered to take the absolute value of the difference from the goal FS ratio, add it to 1, and square the result. Then this value was multiplied by c . For example, if the goal FS ratio was 90 percent, then the fitness could be calculated as

$$\text{MSE ratio} + 400 * (1 + \text{abs}(.90 - (\text{FS}/\text{original FS})))^2 \quad (2)$$

The effect of squaring the difference was to encourage the GA to zero in on the desired FS.

From repeatedly doing runs of 200 population and 500 generations and taking the coefficients and plugging them back into the program as the starting coefficients, runs of over 3,000 generations could be completed. This method introduces slight bias due to the fact that the GA started with the best coefficients from the previous run, but the program randomly perturbs the initial population of best coefficient copies.

3.4.2 Results, Subtask 4

Using the couple image with a quantization step of 64, the enhanced GA written for subtask 4 produced a set of coefficients describing a matched forward and inverse transform pair that produced a 22 percent reduction in MSE in comparison to the performance of Daub4 coefficients and the same FS. Figure 4 shows couple compressed and reconstructed by the Daub4 forward and inverse transform, while Figure 5 shows the same image compressed and reconstructed via the evolved transform pair. Improvements in shading and contrast are visible to the naked eye.



Figure 4. Couple with Quantization Step 64: Daub4 Transform



Figure 5. Couple with Quantization Step 64: Evolved Transform (100 Percent FS)

These coefficients were subsequently tested against each of the images in a test suite. The average over all the images was a 20 percent improvement and the FS was less than 100 percent on average.

Other FS ratios ranging from 90 percent to 110 percent of the wavelet-compressed FS were tried. The results of these tests revealed a trend line that fits the values fairly well. These tests evolved from the same original coefficients (Daub4) and used the couple image each time, so that might have biased the trend line, but there appear to be solutions around the area that correspond to a certain decrease in MSE for a specified FS. Figure 6 shows couple with the FS constraint relaxed to 110 percent of the wavelet-compressed FS. The enhancement in image quality is clearly visible to the naked eye. Figure 7 illustrates the results of evolving matched transform pairs using specific maximum FS values.



Figure 6. Couple with Quantization Step 64: Transform Evolved with Relaxed FS Constraint (110 Percent of Daub4-Compressed FS)




Figure 7. Pareto Optimal Front Illustrating the Tradeoff between FS and MSE Reduction

Next, these evolved coefficients were tested to determine whether they would produce generally good results when used to transform other images. The results of these tests are tabulated in Table 5.

**Table 5. Test Results for Forward and Inverse Transform Pairs
Evolved Using the Couple Image**

	FS (using Daub4)	MSE (using Daub4)	FS (evolved)	MSE (evolved)	FS as % of Daub4 FS	MSE as % of Daub4 MSE	Percentage Improvement of MSE
Airplane	48125	112.4558	47993	88.973	99.726	79.119	20.881
Baboon	48536	279.5612	49758	230.606	102.518	82.488	17.512
Barb	49116	307.6155	49392	257.087	100.562	83.574	16.426
Boat	44429	158.9024	44349	134.718	99.820	84.781	15.219
Couple	44795	155.9632	44792	121.281	99.993	77.763	22.237
Fruits	48640	115.8036	48414	88.458	99.535	76.386	23.614
Goldhill	42814	135.7847	42351	104.912	98.919	77.263	22.737
Lenna	56519	177.5059	56245	137.575	99.515	77.505	22.495
Park	42324	168.9597	42020	133.664	99.282	79.110	20.890
Peppers	39843	121.2649	39683	97.813	99.598	80.661	19.339
Susie	45859	132.5596	45924	106.596	100.142	80.414	19.586
Zelda	38257	134.5674	37903	106.003	99.075	78.773	21.227

Averages 99.89036 79.81961 20.1803875

Coefficients evolved exclusively against couple over approximately 4,000 generations were tested against other images. At quantization 64, they produced an average improvement of 20 percent for MSE while maintaining compressed FS at an average of 99.9 percent when compared to the standard Daub4 transform. There were some instances where FS was slightly larger, the largest being 102.52 percent of the Daub4 compressed file, but on average it appears likely that this will even out. The lowest improvement in MSE over the Daub4 transform was 15 percent against the boat image, while the highest improvement was 23 percent for the fruits image. (Note that this improvement was even greater than for couple.) Thus, the highly evolved coefficients proved quite effective at giving improved MSE when transforming other images at quantization 64 while keeping FS in check.

Figure 8 depicts the fruits image compressed and reconstructed using the Daub4 transforms with a quantization step of 64, while Figure 9 shows the same image processed by the evolved transform pair. Improvements in texture and clarity are visible to the naked eye.



**Figure 8. Fruits Compressed and Reconstructed via the Daub4 Transform:
Quantization = 64**



Figure 9. Fruits Compressed and Reconstructed via Transforms Evolved Against Couple with Identical FS: Quantization = 64

Figure 10 shows fruits after compression and reconstruction using transforms evolved against couple with a quantization step of 64 and a relaxed FS constraint of 110 percent. A considerable amount of additional detail is clearly evident in this image when compared to either the wavelet reconstructed image (Figure 8) or the evolved transform reconstructed image with the 100 percent FS constraint (Figure 9).



Figure 10. Fruits Compressed and Reconstructed by a Transform Evolved Against Fruits and a Relaxed 110 Percent FS Constraint

As a second example of the power of the GA to dramatically improve MSE while increasing FS by only 10 percent, consider Figures 11 and 12. Figure 11 show boats after compression and reconstruction by the Daub4 wavelet with a quantization step of 64, while Figure 12 shows the same image reconstructed via the transform evolved against couple with the same quantization step and a relaxed FS constraint of 110 percent. The latter image reduces MSE in comparison to the former by *a factor of 50.0 percent*.



**Figure 11. Boat Compressed and Reconstructed by the Daub4 Wavelet:
Quantization = 64**



**Figure 12. Boat Compressed and Reconstructed by Evolved Transforms:
Quantization = 64**

Coefficients were also evolved against Baboon, and then Susie, to determine whether the coefficients evolved against other images would also provide good results when applied to other images. Although both of these training runs completed a relatively low number of generations, the results seemed to show good generalization for these coefficients as well. Evolving against Baboon over 500 generations (Table 6) produced coefficients capable of a 6.65 percent improvement in MSE over the Daub4 wavelet at quantization 64. When applied to other test images, these coefficients produced an average improvement of 6.64 percent in MSE. The lowest improvement was 4.9 percent and the highest was 7.4 percent. Evolving against Susie over 500 generations (Table 7) yielded coefficients capable of an 8 percent improvement in MSE over Daub4 at quantization 64. When applied to other test images, these coefficients produced an average improvement of 6.96 percent in MSE. The lowest improvement was 5.4 percent, and the highest was 8.6 percent.

Table 6. Generalization Properties of Coefficients Evolved Using Baboon

	FS (using Daub4)	MSE (using Daub4)	FS (evolved)	MSE (evolved)	FS as % of Daub4 FS	MSE as % of Daub4 MSE	Percentage Improvement of MSE
airplane	48125	112.4558	48099	104.6	99.94597	93.01432	6.985677929
baboon	48536	279.5612	48531	260.97	99.9897	93.34986	6.650135999
Boat	44429	158.9024	44399	150.9	99.93248	94.96395	5.036047284
Fruits	48640	115.8036	48588	107.26	99.89309	92.62234	7.377663561
Lenna	56519	177.5059	56465	164.33	99.90446	92.5772	7.422795524
peppers	39843	121.2649	39803	115.22	99.89961	95.01513	4.984871962
Zelda	38257	134.5674	38211	125.42	99.87976	93.20237	6.797634494

Averages 99.92072 93.53502 6.46497525

Table 7. Generalization Properties of Coefficients Evolved Using Susie

	FS (using Daub4)	MSE (using Daub4)	FS (evolved)	MSE (evolved)	FS as % of Daub4 FS	MSE as % of Daub4 MSE	Percentage Improvement of MSE
Baboon	48536	279.5612	48545	263.59	100.0185	94.28705	5.712953014
Barb	49116	307.6155	49126	291	100.0204	94.59861	5.401385821
Couple	44795	155.9632	44768	144.87	99.93973	92.8873	7.11270351
Goldhill	42814	135.7847	42883	124.09	100.1612	91.38732	8.612678748
Park	42324	168.9597	42394	157.26	100.1654	93.07545	6.924550647
Susie	45859	132.5596	45854	121.95	99.9891	91.99635	8.003645153

Averages 100.049 93.03868 6.961319482

3.4.3 Future Directions

Although most of the test runs for subtask 4 were seeded with the Daub4 coefficients, a run seeded with random coefficients for population 200 over 500 generations (values constrained between -1 and 1) evolved down to 130 percent of MSE and to the same FS as the standard Daub4 transform, leading us to believe it is viable to have more explorative searches of the search space that may find more effective coefficients that are not located so close to the Daub4 coefficients.

Integration of the subimage training methodology of subtask 3 into the software of subtask 4 would result in the creation of a software tool that was capable of much greater MSE reduction for a specified FS, while allowing each training run to complete much more rapidly. Future research will greatly benefit from the combination of these two approaches.

Finally, visible differences between images transformed and reconstructed by wavelets vs. evolved transforms are not as obvious as the MSE measure would seem to indicate. Use of a different fitness measure (such as the Q metric) may allow the GA to produce results that are more apparent to the human eye.

3.5 SUBTASK 5: Evolving New Coefficients for JPEG 2000 Image Compression and Reconstruction

For the fifth subtask identified for this research, the JPEG 2000 Optimizer was used in an attempt to reduce the error incurred during lossy compression of images by the JPEG 2000 compression scheme.

Earl Lamson III completed subtask 5.

3.5.1 Implementation

The Jasper image compression library is used for baseline JPEG 2000 compression. The Jasper library uses the 9/7 wavelet transform. Error reduction is accomplished by optimizing the forward and inverse transforms used during the compression and reconstruction processes. The software developed for subtask 5 uses a GA to simultaneously evolve two sets of coefficients: one set for image compression, and another set for image reconstruction (i.e., each chromosome evolved by the GA contains separate forward and inverse transform coefficients). The application allows the user to customize the optimization process.

3.5.2 GA Details

1. *Generation*

The chromosomes are generated by multiplying each of the coefficients by a random value whose distribution is centered on one and maximum deviation is specified on the command line.

2. *Selection*

Three types of selection operators are included. The selection algorithm to be used is specified on the command line using the 'c' parameter.

- a. Fitness proportionate selection (roulette wheel). This operator is selected by using the command line parameter 'c=0'.
- b. Tournament selection where tournament size is specified on the command line. Tournament selection is specified with 'c=1'. Tournament size is specified by 'T=<tournament size>'.
- c. Linear normalization selection. This operator is specified by using the command line parameter 'c=2'.

3. *Crossover*

The chromosome is represented by an array of doubles encapsulated in a class. Crossover on the chromosomes is performed on the array of doubles and not on the bits. Crossover weight is specified with the 'x' parameter.

4. *Mutation*

Mutation is performed in the same way as generation, except that only one allele (double) is multiplied by the mutation operator. Mutation weight is specified with the 'm' parameter.

5. *Reproduction*

Standard reproduction is used. Reproduction weight is specified with the 'r' parameter.

6. *Evaluation*

Evaluation of a chromosome is done by performing both the forward (compression) and inverse (decompression) transforms on the source image and comparing the results with the original image using the MSE metric. The wavelet coefficients are injected into the Jasper software before both compression and decompression.

3.5.3 Results, Subtask 5

Applying the program to a set of test images yielded the following results:

3.5.3.1 Run 1

Seed: 34

Population Size: 1,000

Generation and Mutation Deviation: 0.01

JPEG 2000 Quality: 0.025 (40:1)

Crossover Weight: 96 percent

Mutation Weight: 3 percent

Reproduction Weight: 1 percent

Selection: Linear Normalization

Source Image: lenna.bmp (512- by 512-pixel color)

Age: 20 generations.

Results:

File	Reference	Evolved	Difference	Percent
lenna.bmp	29.6566	29.5345	0.1221	0.41
airplane.bmp	10.4862	10.6011	-0.1149	-1.10
baboon.bmp	148.977	150.287	-1.31	-0.88
barb.bmp	47.4956	47.1413	0.3543	0.75
boat.bmp	18.1036	18.1843	-0.0807	-0.45
couple.bmp	27.7368	27.8329	-0.0961	-0.35
fruits.bmp	12.2725	12.2995	-0.027	-0.22
goldhill.bmp	27.1412	27.1147	0.0265	0.10
park.bmp	43.7253	43.7158	0.0095	0.02
peppers.bmp	13.8595	13.8335	0.026	0.19
susie.bmp	23.0702	22.8276	0.2426	1.05
zelda.bmp	6.27928	6.26663	0.01265	0.20

Discussion: This test run confirmed the working order of the application. Although progress was occasionally observed, there were not enough generations to draw any real conclusions. Applying the results to the other test images didn't seem to have any useful or predictable effect as it improves some and impairs others.

3.5.3.2 Run 2

Seed: 34

Population Size: 500

Generation and Mutation Deviation: 0.1

JPEG 2000 Quality: 0.0125 (80:1)

Crossover Weight: 90 percent

Mutation Weight: 5 percent

Reproduction Weight: 5 percent

Selection: Linear Normalization

Source Image: lenna_div2.bmp (256- by 256-pixel color)

Age: 300 generations.

Results:

File	Reference	Evolved	Difference	Percent
lenna_div2.bmp	122.196	119.908	2.288	1.87
airplane_div2.bmp	118.986	111.079	7.907	6.65
baboon_div2.bmp	269.712	259.828	9.884	3.66
barb_div2.bmp	143.031	144.12	-1.089	-0.76
boat_div2.bmp	114.149	114.988	-0.839	-0.74
couple_div2.bmp	131.037	131.795	-0.758	-0.58
fruits_div2.bmp	72.7055	72.1221	0.5834	0.80
goldhill_div2.bmp	85.3961	82.7413	2.6548	3.11
park_div2.bmp	200.942	207.036	-6.094	-3.03
peppers_div2.bmp	70.449	71.8374	-1.3884	-1.97
susie_div2.bmp	76.9911	73.1166	3.8745	5.03
zelda_div2.bmp	26.7705	26.2331	0.5374	2.01

To attempt to compensate for the slow speed of the computer on which this was run, the resolution of the test images was reduced. More improvement is made on this run over the first, as is to be expected with the greater number of generations. The improvement is still too insignificant to be visible. Applying the evolved coefficients to the rest of the images again tends to have unpredictable results.

3.5.3.3 Run 3

Seed: 34

Population Size: 1,000

Generation and Mutation Deviation: 0.01

JPEG 2000 Quality: 0.025 (40:1)

Crossover Weight: 96 percent

Mutation Weight: 3 percent

Reproduction Weight: 1 percent

Selection: Linear Normalization

Source Image: goldhill.bmp (512- by 512-pixel grey)

Age: 325 generations.

Results:

File	Reference	Evolved	Difference	Percent
goldhill.bmp	27.1412	26.7477	0.3935	1.45
airplane.bmp	10.4862	10.7184	-0.2322	-2.21
baboon.bmp	148.977	151.117	-2.14	-1.44
barb.bmp	47.4956	47.0334	0.4622	0.97
boat.bmp	18.1036	17.9019	0.2017	1.11
couple.bmp	27.7368	27.9996	-0.2628	-0.95
fruits.bmp	12.2725	12.4802	-0.2077	-1.69
lenna.bmp	29.6566	29.5414	0.1152	0.39
park.bmp	43.7253	43.5204	0.2049	0.47
peppers.bmp	13.9767	13.9409	0.0358	0.26
susie.bmp	23.0702	22.8128	0.2574	1.12
zelda.bmp	6.27928	6.27881	0.00047	0.01

This run was the longest of the four, although it remains quite young in the scope of a GA. The decreased frequency of improvement in later generations discouraged spending more time aging this run. Unfortunately, very small improvement of the already small error metric was observed. The initially high quality images led to the change in parameters for the last run.

3.5.3.4 Run 4

Seed: 34

Population Size: 500

Generation and Mutation Deviation: 0.01

JPEG 2000 Quality: 0.0125 (80:1)

Crossover Weight: 90 percent

Mutation Weight: 5 percent

Reproduction Weight: 5 percent

Selection: Linear Normalization

Source Image: goldhill.bmp (512- by 512-pixel grey)

Age: 250 generations.

Results:

File	Reference	Evolved	Difference	Percent
goldhill.bmp	52.2814	51.338	0.9434	1.80
airplane.bmp	29.2861	28.8681	0.418	1.43
baboon.bmp	285.045	285.939	-0.894	-0.31
barb.bmp	105.492	105.725	-0.233	-0.22
boat.bmp	45.2636	44.4933	0.7703	1.70
couple.bmp	63.7842	63.8161	-0.0319	-0.05
fruits.bmp	26.46	26.8051	-0.3451	-1.30
lenna.bmp	49.2191	49.2934	-0.0743	-0.15
park.bmp	84.2667	83.7834	0.4833	0.57
peppers.bmp	24.3578	24.4653	-0.1075	-0.44
susie.bmp	46.2359	45.8857	0.3502	0.76
zelda.bmp	10.1934	10.3685	-0.1751	-1.72

This last run was an attempt to improve upon the results of run 3 by increasing the compression rate of the Jasper library by a factor of 2. The increase compression does indeed lead to higher error levels; unfortunately, the improvement of this higher error level is again insignificant. The population size was reduced due to the appearance of many copies of the same chromosome in the generations of run 3. The changes don't seem to help. These results indicate that the performance improvement over the reference wavelet on images remains chaotic. With both runs 3 and 4 being trained on the same image, it was hoped that the evolved coefficients would have similar effects when applied to the other images, but this does not appear to be the case.

3.5.4 Conclusions, Subtask 5

The results of subtask 5 were inconclusive. Images produced the transforms evolved during this subtask do not appear to have improved upon the JPEG 2000 standard in any visible sense. From a purely numerical perspective, the GA was occasionally successful in reducing the error introduced by compression, although the metric used for measuring the error may not reflect the amount of change made by the evolved coefficients.

3.5.5 Continuation

Future research should concentrate on the following improvements:

1. Error metrics

Incorporate other error metrics, as well as the ability to select among them. Also analysis of the error in the images would be nice, for instance a histogram showing the distribution of per pixel error values. It is suspected that improvements are being made by reducing a few large error values in favor of many smaller error values.

2. Fitness hash optimization

Much redundancy in the population causes the same chromosome to be repeatedly evaluated. Instead of straight evaluation of each chromosome, the fitness for each unique chromosome encountered could be stored in a hash table. The hash table could then be checked for duplicates before evaluation, saving a possibly very large amount of time for larger populations. Of course, this method would cost the memory of storing this table of fitness values, but the memory footprint of the current program is fairly small in comparison to the memory available in today's computers.

3. *User interface*

Either a graphical user interface (GUI) or a more robust text user interface would make working with the runs much easier. This change would require a redesign of some of the object interfaces. Of course, this re-factoring probably needs to be done regardless, as many afterthought features have revealed weakness in the application's design.

4. *Specialized reproduction (elitism)*

Elitism would keep the run's best chromosome in the gene pool and could increase the rate of convergence.

5. *Platform independence*

Originally, the application was designed to run on *nix machines in a clustered or massively parallel computing environment. Although the current code does compile and run on *nix machines, segmentation faults frequently occur after a few generations, likely due to a memory issue. Future efforts should focus upon debugging the application.

4. CONCLUSIONS

This research project resulted in several important contributions to the state of the art in the image processing field.

- A. This study demonstrated that a GA is capable of evolving coefficients representing inverse transforms that outperform wavelets on the task of minimizing the MSE in reconstructed images under conditions subject to quantization.
- B. The results of subtask 1 showed that inverse transforms evolved using a single representative image also outperformed wavelets when subsequently tested against other images.
- C. Subtask 2 established a methodology for using one or more representative subimages to evolve coefficients representing optimized inverse transforms, and showed that the evolved transforms also outperformed wavelets in terms of MSE reduction in reconstructed images. The results summarized above demonstrated that the use of representative subimages to evolve inverse transform coefficients greatly increased the speed of the GA, allowing the use of much larger values for GA control parameters M (population size) and G (number of generations executed).
- D. Subtask 2 also investigated several approaches to the problem of evolving an inverse transform capable of highlighting the existence of a selected subimage in a larger parent image. One approach that was attempted was to create a training population that consisted of both the selected subimage and the parent image, and then utilize a fitness function that simultaneously encouraged high-fidelity reconstruction of the subimage and poor reconstruction of the parent image. Unfortunately, the transforms developed during subtask 2 were incapable of using simple multiobjective fitness criteria to solve this problem.
- E. Subtask 3 attempted to discover the properties of subimages that make them more or less suitable for evolving an inverse transform capable of performing well over all of the test images. A representativeness measure was generally successful in predicting which subimages could be used to evolve a better inverse transform. In addition, the results of subtask 3 suggested that other factors, such as the clarity and texture of the subimage, may be a more reliable measure for selecting the best subimage for training.
- F. During subtask 4, the previously developed model was extended to allow the simultaneous evolution of coefficients representing matched transform and inverse transform pairs. This enhancement over the previous model (which supported only the evolution of inverse transform coefficients) resulted in a much higher degree of MSE reduction. According to the MSE metric, the quality of images transformed and subsequently reconstructed by evolved transform pairs was significantly better than the quality of images transformed by a wavelet and subsequently reconstructed by an evolved inverse transform, which in turn was better than the quality of images transformed and reconstructed using standard wavelets under conditions subject to quantization error.

5. RECOMMENDATIONS FOR FUTURE RESEARCH

Based upon the outcome of this research project, the PI recommends future research to be conducted into the following issues:

- A. The outcome of subtask 4 clearly demonstrated the increased power of simultaneously evolving matched forward and inverse transform pairs for high-fidelity image compression and reconstruction under conditions subject to quantization error. To dramatically increase both the speed of the evolutionary process and the size of runs that could be completed, the subimage training techniques identified during subtask 2 should be integrated with the approach taken during subtask 4. The resulting software may be capable of achieving far more impressive results than has been observed to date.
- B. The outcome of subtask 3 demonstrated the need to investigate the properties of subimages that make them more or less useful to the GA for evolving optimized transform coefficients. The resulting subimage selection criterion could then be integrated directly into the GA, allowing it to automatically identify the most suitable subimages for training.
- C. It is possible that the integration of the above methodologies will result in a GA capable to evolving better coefficients than are currently utilized by the JPEG 2000 standard.
- D. The use of other fitness measures besides MSE (such as the Q metric) may allow the GA to evolve forward and inverse transform pairs that produce reconstructed images that are much more appealing to the human eye than images produced by standard wavelets.
- E. To date, generation 0 used by each GA run has been populated with randomly mutated copies of coefficients describing a specified standard wavelet. Since mutation occurs infrequently and consists of multiplying a given coefficient by a factor chosen from a narrow Gaussian distribution around 1.0, virtually all of the individuals in generation 0 of each run existed in the transform space immediately adjacent to the wavelet. Furthermore, all of the transforms evolved by the GA developed for this study have utilized a structure that is identical to a selected wavelet. (For most of the examples described in this report, the Daub4 wavelet was used to seed the initial population of transforms; thus, all of the forward or inverse transforms from each generation of each GA run were comprised of exactly four g and four h coefficients.) However, a growing amount of empirical evidence suggests that transforms radically different from wavelets in both structure and composition may be better at image compression and reconstruction, offsetting detrimental effects such as quantization more effectively. Future research should investigate the use of populations of evolved transforms that vary in both shape and composition.
- F. The full impact of this research will not be appreciated by the signal and image processing community until its successful application to a real-world system. Future efforts should focus upon the identification of systems that will benefit the most from this GA-based approach.

6. REFERENCES

- Astrom, K. J., G. C. Goodwin, and P. R. Kumar 1993. *Adaptive Control, Filtering, and Signal Processing*, Institute for Mathematics and its Applications Vol. 74, ISBN 0-387-97988-3.
- Bradley, J. N., C. M. Brislawn, and T. Hopper 1993. The FBI Wavelet-Scalar Quantization Standard for Grey-Scale Fingerprint Compression, *Visual Information Processing II*, Vol. 1961: 293-304, *Proceedings of the SPIE*, Orlando, FL.
- Brartke, S. J. and A. Barto 1996. Linear Least-Squares Algorithms for Temporal Difference Learning, *Machine Learning*, 22: 33-57.
- Bruckmann, A., T. Schell, and A. Uhl 1998. Evolving Subband Structures for Wavelet Packet Based Image Compression Using Genetic Algorithms with Non-Additive Cost Functions, *Proceedings of the International Conference on Wavelets and Multiscale Methods (IWC'98)*, INRIA.
- Burrus, C. S., R. A. Gopinath, and H. Guo 1998. *Introduction to Wavelets and Wavelet Transforms: A Primer*, Prentice-Hall.
- Candocia, F. M. and J. C. Principe 2000. Superresolution of Images with Learned Multiple Reconstruction Kernels, *Multimedia Image and Video Processing*: 67-95, CRC Press.
- Chang, S. 1995. *Image Interpolation Using Wavelet-Based Edge Enhancement and Texture Analysis*, Master's Thesis, University of California at Berkeley.
- Chang, G., B. Yu, and M. Vetterli 1998. Spatially adaptive wavelet thresholding based on context modeling for image denoising, *IEEE Transactions on Image Processing*.
- Charrier, M., D. S. Cruz, and M. Larsson 1999. JPEG2000, the Next Millennium Compression Standard for Still Images, *IEEE International Conference on Multimedia Computing and Systems*, 1: 131-132.
- Claypoole, R. L., R. G. Baraniuk, and R. D. Nowak 1999. Adaptive Wavelet Transforms via Lifting, Technical Report 9304, Rice University.
- Cohen, A., I. Daubechies, and J. C. Feauveau, 1992. Biorthogonal Bases of Compactly Supported Wavelets, *Communications on Pure and Applied Mathematics*, 45: 485-560.
- Cooley, J. W. and J. W. Tukey 1965. An Algorithm for the machine calculation of complex Fourier series, *Mathematics of Computation*, 19: 297-301.
- Cornall, T. *et al.* 1998. JPEG Image Coding Standard. The Centre for Telecommunications and Information Engineering (CTIE). Accessed March 25th, 2005.
<http://www.ctie.monash.edu.au/EMERGE/multimedia/JPEG/>.
- Daubechies, I. 1988. Orthonormal Bases of Compactly Supported Wavelets, *Communications of Pure and Applied Mathematics*, 41:909-996.
- Daubechies, I. 1992. *Ten Lectures on Wavelets*, SIAM.

- Davis, G. M. and A. Nosratinia 1998. Wavelet-based Image Coding: An Overview, *Applied and Computational Control, Signals, and Circuits*, 1(1).
- Donoho, D. 1993. Nonlinear Wavelet Methods for Recovery of Signals, Densities, and Spectra from Indirect and Noisy Data, in Daubechies, I. (Ed.), *Different Perspectives on Wavelets*, AMS.
- Donoho, D. and I. Johnstone 1994. Ideal Spatial Adaptation via Wavelet Shrinkage, *Biometrika*, 81: 425-455.
- Downie, T. R. 1998. Signal Preprocessing of Multiwavelets. Available from <http://citeseer.ist.psu.edu/downie98signal.html>.
- Fante, R.L., 1988. *Signal Analysis and Estimation: an Introduction*, John Wiley.
- Farhang-Boroujeny, B. 1999. *Adaptive Filters: Theory and Applications*, John Wiley.
- Garza, A. G. and M. L. Maher 1999. An Evolutionary Approach to Case Adaptation, *Third International Conference on Case-Based Reasoning (ICCBR'99)*, Springer-Verlag.
- Gersho, A. and R. M. Gray 1991. *Vector Quantization and Signal Compression*, Kluwer Academic Publishers.
- Goldberg, D. E. 1989. *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley.
- Graps, A. 1995. An Introduction to Wavelets, *IEEE Computational Science and Engineering*, 2(2), IEEE.
- Haar, A. 1910. Zur Theorie der orthogonalen Funktionensysteme, *Mathematische Annalen*, 69: 331-371.
- Haykin, S. 1996. *Adaptive Filter Theory* (3rd edition), Prentice Hall.
- Haykin, S. and B. Widrow (Eds.) 2003. *Least-Mean-Square Adaptive Filters*, John Wiley.
- Holland, J. 1975. *Adaptation in Natural and Artificial Systems*, University of Michigan Press.
- Jones, E., P. Runkle, N. Dasgupta, and L. Carin 2000. Signal Adaptive Wavelet Design using GA, *Proceedings of the SPIE*, 4056: 362-371.
- Klappenecker, A. and F. U. May 1995. Evolving better wavelet compression schemes, *Wavelet Applications in Signal and Image Processing III*, SPIE Proceedings, 2569.
- Kobayashi, K. and T. Torioka 1997. Designing Wavelet Networks Using Genetic Algorithms, *Proceedings of 5th European Congress on Intelligent Techniques and Soft Computing (EUFIT'97)*, 1: 429-433, Verlag Mainz.
- Koza, J. R. 1992. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*, MIT Press.

- Lai, Y.-K. and C.-C. J. Kuo, 1997. Image quality measurement using the Haar wavelet, *Wavelet Applications in Signal and Image Processing V*, SPIE.
- Langdon, W. B. and J. P. Nordin 2000. Seeding Genetic Programming Populations, *Lecture Notes in Computer Science*, 1802: 304-315, Springer-Verlag.
- Lankhorst, M. M. and M. D. van der Laan 1995. Wavelet-based Signal Approximations with Genetic Algorithms, *Proceedings of the 4th Annual Conference on Evolutionary Programming*, 237-255, MIT Press.
- Le Pennec, E. and S. Mallat 2000. Image Compression with Geometrical Wavelets, *IEEE Conference on Image Processing*, IEEE Signal Processing Society.
- Lin, C. 1999. *Time and Space Efficient Wavelet Transform for Real-Time Applications*, doctoral dissertation, The Ohio State University, 161-2.
- Long, C. J. and S. Datta 1998. Discriminant Wavelet Basis Construction for Speech Recognition, *5th International Conference on Spoken Language Processing (ICSLP 98)*, 3: 1047-1050.
- Malfait, M. 1996. Using Wavelets to Suppress Noise in Biomedical Images, in Aldroubi, A. and M. Unser (Eds.), *Wavelets in Medicine and Biology*, CRC Press.
- Mallat, S. 1989. A Theory for Multiresolution Signal Decomposition: The Wavelet Decomposition, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11: 674-693.
- Mallat, S. 1998. *A Wavelet Tour of Signal Processing*, Academic Press.
- Miller, B. L. and D. E. Goldberg 1995. Genetic Algorithms, Tournament Selection, and the Effects of Noise, *Complex Systems*, 9:193-212.
- Murray-Smith, D. J., E. W. McGookin, T. I. Fossen, and Y. Li 2000. Ship Steering Control System Optimization Using Genetic Algorithms, *Control Engineering Practice*, 8: 429-443, Elsevier Science.
- Nag, S. K. and L. Peters, Jr. 1998. Ramp Response Signatures for Dielectric Targets, *SPIE Proceedings*, 3392: 703-713, SPIE.
- Nitzberg, M. and T. Shiota 1992. Nonlinear Image Filtering with Edge and Corner Enhancement, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14 (8): 826-833, IEEE.
- Noel, S. and H. Szu 1998. Wavelets and Neural Networks for Radar, *Progress in Unsupervised Learning of Artificial Neural Networks and Real-World Applications*, Russian Academy of Nonlinear Sciences, ISBN 0-620-23629-9.
- Odegard, J., R. Gopinath, and C. Burrus 1994. Design of Linear Phase Cosine Modulated Filter Banks for Subband Image Compression, CML Technical Report, Rice University.
- Piella, G. and H. J. A. M. Heijmans 2001. Adaptive Lifting Schemes with Perfect Reconstruction, Report PNA-R0104, Centrum voor Wiskunde en Informatica, Stichting Mathematisch Centrum.

- Preston, K. Jr. 1992. Three-dimensional Morphology Based on Z-filters, *Nonlinear Image Processing III*, SPIE Proceedings 1658: 63-74, SPIE.
- Rajoub, B. A. 2002. An Efficient Coding Algorithm for the Compression of ECG Signals Using the Wavelet Transform, *IEEE Transactions on Biomedical Engineering*, 49(4): 355-362.
- Saha, S. 2000. Image Compression – from DCT to Wavelets: A Review, *ACM Crossroads*, ACM.
- Saha, S. and R. Vemuri 1999. Adaptive Wavelet Filters in Image Coders – How Important Are They?, *IECON'99 Proceedings: The 25th Annual Conference of the IEEE Industrial Electronics Society*, 2: 559-564, IEEE Industrial Electronics Society.
- Saha, S. and R. Vemuri 2000. Analysis-based Adaptive Wavelet Filter Selection in Lossy Image Coding Schemes, *Proceedings: IEEE International Symposium on Circuits and Systems (ISCAS 2000)*, Geneva, Switzerland.
- Saito, N. 1994. Simultaneous Noise Suppression and Signal Compression using a Library of Orthonormal Bases and the Minimum Description Length Criterion, in Foufoula-Georgiou, E. and P. Kumar (Eds.), *Wavelets in Geophysics, Wavelet Analysis and Its Applications Vol. 4*, Academic Press.
- Savale, B., O. Delage, and H. H. Arsenault 1998. Iterative denoising and simultaneously reconstructing algorithm: a new approach for images degraded by multiplicative signal-dependent noise, *SPIE Proceedings: Applications of Digital Image Processing XXI*, 3460: 811-820, SPIE.
- Schilling, D., P. Cosman, and C. Berry 1998. Image Recognition in Single-scale and Multiscale Decoders, *Proceedings of the 32nd Asilomar Conference on Signals, Systems, and Computers*, IEEE Computer Society Press.
- Schremmer, C., T. Haenselmann, and F. Bömers 2001. A Wavelet Based Audio Denoiser, *IEEE International Conference on Multimedia and Expo (ICME) 2001*, 145-148.
- Shannon, C. E. and W. Weaver 1964. *The Mathematical Theory of Communication*, University of Illinois Press.
- Skodras, A., C. Christopoulos, and T. Ebrahimi 2001. The JPEG 2000 Still Image Compression Standard, *IEEE Signal Processing Magazine*, Sept. 2001: 37-58.
- Stein, E. M. and G. Weiss 1971. *Introduction to Fourier Analysis in Euclidean Spaces*, Princeton Mathematical Series 32, Princeton University Press.
- Takehisa, Y., H. Sakanashi, and T. Higuchi 2000. Adaptive Wavelet Transform for Lossless Compression using Genetic Algorithm, *GECCO-2000: Proceedings of the Genetic and Evolutionary Computation Conference, July 10-12, 2000, Las Vega, NV*, Morgan Kaufmann.
- Tinsley, D. and J. Kettell 2004. A Genetic Algorithm for Subimage Optimization using Wavelets, Final Report, CS A490 Genetic Algorithms, University of Alaska Anchorage, Fall 2004.

- Usevitch, B. E. 2001. A Tutorial on Modern Lossy Wavelet Image Compression: Foundations of JPEG 2000, *IEEE Signal Processing Magazine*, Sept. 2001: 22-35.
- Vaidyanathan, P. P. 1992. *Multirate Systems and Filter Banks*, Prentice-Hall.
- Venkatachalam, V. and J. Aravena 1998. Detecting Periodic Behavior in Nonstationary Signals, *IEEE-SP International Symposium on Time-frequency and Time-scale Analysis*, IEEE.
- Vetterli, M. and J. Kovacevic 1995. *Wavelets and Subband Coding*, Prentice-Hall.
- Villasenor, J., B. Belzer, and J. Liao 1995. Wavelet filter evaluation for image compression, *IEEE Transactions on Image Processing*, 4(8): 1053-1060.
- Walker, J. S. 1999. *A Primer on Wavelets and their Scientific Applications*, CRC Press.
- Wang, Y., Q. Wu, K. R. Castleman, and Z. Xiong 2001. Image Enhancement Using Multiscale Oriented Wavelets, *Proceedings, IEEE International Conference on Image Processing (ICIP'01)*, Thessalonica, Greece.
- Zhu, Y. and S. Schwartz 2002. Discriminant Analysis and Adaptive Wavelet Feature Selection for Statistical Object Detection, *Proceedings, 16th International Conference on Pattern Recognition (ICPR'02)*, 4: 40086-40089, IEEE.